

Uniform Information Exchange in Multi-channel Wireless Ad Hoc Networks

Li Ning

Center for High Performance Computing
Shenzhen Institutes of Advanced Technology, CAS
Shenzhen, P.R. China
li.ning@siat.ac.cn

Yong Zhang

Center for High Performance Computing
Shenzhen Institutes of Advanced Technology, CAS
Shenzhen, P.R. China
zhangyong@siat.ac.cn

Francis C.M. Lau

Department of Computer Science
The University of Hong Kong
Hong Kong, P.R. China
fcmlau@cs.hku.hk

Dongxiao Yu

Department of Computer Science
The University of Hong Kong
Hong Kong, P.R. China
dxyu@cs.hku.hk

Yuexuan Wang

Department of Computer Science
The University of Hong Kong
Hong Kong, P.R. China
amywang@hku.hk

Shengzhong Feng

Center for High Performance Computing
Shenzhen Institutes of Advanced Technology, CAS
Shenzhen, P.R. China
sz.feng@siat.ac.cn

November 19, 2016

Abstract

In the information exchange problem, k packets that are initially maintained by k nodes need to be disseminated to the whole network as quickly as possible. We consider this problem in single-hop multi-channel networks of n nodes, and propose a uniform protocol that with high probability accomplishes the dissemination in $O(k/\mathcal{F} + \mathcal{F} \cdot \log n)$ rounds, assuming \mathcal{F} available channels and collision detection. This result is asymptotically optimal when k is large ($k \geq \mathcal{F}^2 \cdot \log n$). To our knowledge, this is the first uniform protocol for information exchange in multi-channel networks.

1 Introduction

In this paper, we study the information exchange problem in a single-hop, multi-channel radio network. There are k nodes, called the *source nodes*, in the network. At the beginning, each of them holds a packet, and the target is to disseminate these packets to the whole network as quickly as possible. Information exchange is one of the most fundamental operations that are frequently called for in the smooth running of a network.

Using multiple channels obviously can greatly increase the throughput of the network. A lot of works have been devoted to studying the utilization of multiple channels in the derivation of faster communication protocols (e.g. [5, 8, 6, 9, 10, 12, 11, 14, 15, 19, 20]). All existing works however require that the network size n be known a priori. In ad hoc networks, knowing n is usually a tough task, as it would consume a large amount of time and energy for nodes to compute this global parameter, and hence greatly increase the load of the network. Additionally, in ad hoc networks, the network size could change frequently due to nodes leaving and joining. This consideration necessitates the design of uniform protocols which do not require any prior information about network parameters including the network size n and the number of source nodes k . Uniform protocols have better scalability and therefore more suitable for implementation in reality. In this paper, we propose a uniform protocol for information exchange whose time complexity decreases linearly as the number of available channels increases.

1.1 Network Model and Problem Definition

A multi-channel single-hop network is defined as follows. There are n nodes in the network, and any pair of which can communicate with each other directly. But n is not known to nodes. Time is divided into synchronous rounds. There are \mathcal{F} channels available in the network. We use $1, \dots, \mathcal{F}$ to denote these channels. Even though these \mathcal{F} channels are available to all the nodes, at any time a node can select at most one channel to listen to or transmit on. A node operating on a channel in a given round learns nothing about events on the other channels. When a node v listens to a channel, it can receive a message if and only if there is only one node transmitting on the channel. If two or more nodes transmit on the same channel, a collision occurs and none of these transmissions would be successful. We assume that nodes can detect collisions, i.e., nodes can distinguish collision from silence. Furthermore, we consider the case of non-constant \mathcal{F} (larger than any constant), since otherwise, using a constant number of channels will not break the $\Omega(k)$ lower bound for information exchange that always holds in single-channel networks.

The algorithm proposed in this paper is randomized, and hence the analysis involves many random events. We say that an event happens with high probability (with respect to n), if it happens with probability $1 - 1/n^c$ for some constant $c > 0$.

The goal of information exchange is to disseminate some source nodes' packets to the whole network, which is more precisely defined as follows.

Definition 1. (Information Exchange.) *In the information exchange problem, initially k source nodes are holding packets $\{P_1, P_2, \dots, P_k\}$ respectively. It is required to disseminate all these k packets to the whole network as quickly as possible.*

Denote by K the set of source nodes. Then $|K| = k$. We study the harsh case of the information exchange problem where nodes have no idea about the number of packets k and the set of source nodes K . We assume that multiple packets can be packed in a single message. It is easy to see that if k is small relative to the number of channels \mathcal{F} , the benefit of multiple channels will be weakened, since in this case there could be a single node selecting a channel such that its transmission cannot be received by anyone. Thus, throughout this work, we assume that $k \geq \mathcal{F} \log n$, which ensures that when nodes uniformly select the channels, there are multiple nodes operating on each channel with high probability. However, we must point out that our algorithm can also solve the case where k is small.

1.2 Our Result and Technique

In this paper, we give the first known uniform protocol for information exchange in multi-channel networks. Our algorithm can disseminate all k packets to the whole network in $O(k/\mathcal{F} + \mathcal{F} \cdot \log n)$ rounds with high probability when there are \mathcal{F} available channels. When k is large ($k \geq \mathcal{F}^2 \log n$), our algorithm shows a linear speedup considering the $\Omega(k)$ lower bound for single-channel networks. Note that $\Omega(k/\mathcal{F})$ is a trivial lower bound for information exchange with \mathcal{F} available channels. Hence, our protocol is asymptotically optimal when k is large.

In our protocol, every node that needs to transmit maintains a transmission probability, and in each round, a node decides to transmit with its transmission probability. With \mathcal{F} available channels, our protocol applies a very intuitive rule for the nodes to do the selection: in each round, a node just selects one channel uniformly at random, and then transmits or listens on the selected channel. If a node listens on the channel and detected that the selected channel is *idle*, then it doubles its transmission probability. Otherwise, the node halves the transmission probability.

By a straightforward computation, it is easy to discover that in order to ensure a successful transmission on one channel with constant probability, the total transmission probability of nodes selecting this channel should be a constant. Hence, to efficiently make use of the channels, the total transmission probability of nodes should be in a "safe range" $[\alpha_1 \cdot \mathcal{F}, \alpha_2 \cdot \mathcal{F}]$ with constants $\alpha_1, \alpha_2 > 0$. However, this is not easy to achieve without a careful design of the protocol. The difficulty comes from the selection process of the channels. Since the nodes select the channels using a distributed, randomized protocol and the selections of nodes are mutually independent, the total transmission probability of nodes selecting a channel may vary a lot among different channels. As a result, nodes may update their transmission probabilities towards different directions, which makes it very hard to analyze whether the "safe range" is still guaranteed after an update. Our protocol has a channel-consistent-updating property, i.e., nodes selecting the same channel update their transmission probabilities consistently. With this channel-consistent-updating property, our analysis shows that when the total transmission probability of all the nodes goes outside the "safe range", there are enough channels on which

the nodes behave consistently to pull the total transmission probability back to the safe range. Then as a global effect, the network stays stable in the “safe range” state, and the \mathcal{F} available channels are used efficiently (i.e. transmit successfully $\Omega(\mathcal{F})$ messages with constant probability in each round).

Notice that if we can quickly aggregate all the k packets to a single node and let this node be the only one to transmit in the network (by broadcasting in a pre-defined primary channel), then the whole network will know all the k packets very soon. Our protocol follows this approach, and hence prevents the node from trying to transmit anymore if its message has been successfully received by some node which is still transmitting. This technique is also known as *indirecting*, and can be summarized as “if your message is received by another speaker, then you never speak again”. By this approach, the nodes that try to transmit become fewer and fewer as the protocol is running.

When the number of transmitting nodes becomes small, using all $\mathcal{F} > 1$ channels does not always help. The reason is when there are only a few nodes to transmit, it is hard for them to meet each other if they still select the channel randomly. A direct solution will be such that if the transmitting nodes find that there are only a few left nodes, then they stop selecting channels and operate on a pre-defined channel. The primary channel that is designed for the final broadcast can be used to achieve this. However, the problem has not been completely solved. By our analysis, we know that the number of rounds needed before the multiple channels become inefficient is $T = O(\log n + k/\mathcal{F})$ with high probability. Unfortunately, this time bound cannot be known to the nodes in the network, since the protocol is uniform and information about n and k is not available to the nodes. Consequently, the nodes cannot calculate T and have no idea about when to stop selecting channels. Our algorithm uses a parallel approach to overcome this difficulty. In our protocol, each round is divided into four slots: in the first two slots, nodes use multiple channels for transmissions; and in the last two slots, all nodes only use the pre-defined primary channel for transmission. This parallel approach affects the running time by only a constant factor, but perfectly solves the inefficiency problem of multiple channels for information exchange with small number of transmitting nodes. Our analysis shows that such a parallel approach completes the information exchange in $O(k/\mathcal{F} + \mathcal{F} \cdot \log n)$ rounds with high probability.

1.3 Related Work

As more and more wireless networks and devices now operate on multiple channels, there has been much attention given to studying the effect of multiple channels on facilitating communication recently [5, 8, 6, 9, 10, 12, 11, 14, 15, 19, 20]. With respect to information exchange in multi-channel single-hop networks, most studies are done under the assumption that each message can carry only one packet. In particular, Holzer et al. [15, 14] proposed deterministic and randomized algorithms with optimal $O(k)$ time to solve the information exchange problem. With the assumption that nodes can listen to and receive messages from multiple channels at the same time, Shi et al. [19] proposed an $O(\log k \log \log k)$ time randomized information exchange protocol using $\Theta(n)$ channels. But with the assumption of unit-size messages, the benefit of utilizing multiple channels is very limited, since in each round, a node can receive at most one packet. Hence, it needs $\Omega(k)$ rounds to complete the information exchange. On the other hand, the packet stored at nodes could be small (e.g., in sensor networks, the data at each node is only a value). It is realistic to consider the case that multiple packets can be packed in a single message. Under this assumption, in [6], Daum et al. proposed a randomized algorithm that accomplishes information exchange in $O(k + \log^2 n/\mathcal{F} + \log n \log \log n)$ rounds with high probability. Their algorithm does not rely on collision detection. Then with collision detection, Wang et al. [20] proposed a protocol that disseminates all the packets in $O(k/\mathcal{F} + \mathcal{F} \cdot \log^2 n)$ rounds with high probability. When k is large ($k \geq \mathcal{F}^2 \log^2 n$), this result is asymptotically optimal considering the trivial lower bound $\Omega(k/\mathcal{F})$. In [22], Yan et al. studied the impact of message size on information exchange in multi-channel networks. Additionally, Gilbert et al. [12] considered the scenario when an adversary can disrupt a number of channels and proposed a randomized algorithm to achieve the almost-complete information exchange. However, all the above results need the prior knowledge of n . To our knowledge, there is not yet any uniform protocol proposed for solving the information exchange problem in single-hop multiple-channel networks.

Information exchange has also been extensively studied since 1970s [4, 13, 18] in single-channel networks. In single-channel networks, information exchange is also known as contention resolution [2] or k -selection [16]. Assuming collision detection as in this work, a randomized adaptive protocol with expected running time of $O(k + \log n)$ was presented by Martel in [17]. Kowalski [16] improved the protocol in [17] to $O(k + \log \log n)$ by making use of the expected $O(\log \log n)$ selection protocol in [21]. When requiring high probability results, the best known randomized algorithm was introduced in [1], which solves the k -selection problem in $O(k + \log^2 n)$ rounds without assuming collision detection. Note that in the single-channel networks, the trivial lower bound

for k -selection is $\Omega(k)$. Hence the result in [1] is asymptotically optimal for $k \in \Omega(\log^2 n)$. By assuming that the channel can provide feedback on whether a message is successfully transmitted, a uniform randomized protocol with running time $O(k)$ is introduced in [2] for single-channel networks. However, the error probability of the protocol in [2] is $1/k^c$, rather than $1/n^c$. For deterministic solutions, adaptive protocols for k -selection were presented with running time $O(k \log(n/k))$ in [4, 13, 18], assuming collision detection.

1.4 Outline

Section 2 introduces some preliminary results that help the analysis. Section 3 introduces our protocol. Section 4 analyzes the performance of our protocol; particularly, we give an upper bound on the time needed to accomplish (with high probability) information exchange. Furthermore, we show the “self-stabilization” property of our protocol. Section 5 summarizes our work, followed by a discussion.

2 Preliminaries

In this section, we review some useful results concerning randomness.

Lemma 1 (Chernoff Bound.). *Consider a set of random variables $0 \leq X_1, X_2, \dots, X_n \leq c$ for some parameter $c > 0$. Let $X := \sum_{i=1}^n X_i$ and $\mu := \mathbb{E}[X]$. If X_i ’s are independent or negatively associated, then for any $\delta > 0$ it holds that*

$$\Pr[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\frac{\mu}{c}}.$$

In details, for $\delta \leq 1$, the bound can be upper bounded by

$$\Pr[X \geq (1 + \delta)\mu] \leq \exp\left\{-\frac{\delta^2 \mu}{3c}\right\};$$

for $\delta > 1$, it holds that

$$\Pr[X \geq (1 + \delta)\mu] \leq \exp\left\{-\frac{\delta \ln(1 + \delta)\mu}{2c}\right\}.$$

On the other hand, for any $0 < \delta < 1$ it holds that

$$\Pr[X \leq (1 - \delta)\mu] \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^{\frac{\mu}{c}} \leq \exp\left\{-\frac{\delta^2 \mu}{c}\right\}.$$

Next, we present some useful conclusions about a classic procedure, “throw balls into bins”. These conclusions have essentially been proved in existing works such as [7]. However, for the completeness of our arguments, we will go through the proof in details.

Lemma 2. *Consider H bins and l balls with weights $0 \leq w_1, w_2, \dots, w_l \leq \zeta$. Assume that $\sum_{i=1}^l w_i = \alpha \cdot H$ where $\alpha \geq 0.01$ is a constant. Balls are thrown into bins uniformly at random. Then, if ζ is small enough, with probability $1 - \exp\{-\Omega(H)\}$ there are at least $H \cdot 31/32$ bins in which the total weight of balls is between $\alpha \cdot 15/16$ and $\alpha \cdot 2$.*

Proof. Considering the j -th bin, let X_j^i denote the random variable that takes value w_i if the i -th ball is in the j -th bin, and 0 otherwise. Since the balls are thrown to bins uniformly at random, $\Pr[X_j^i = w_i] = 1/H$ and $\mathbb{E}[\sum_i X_j^i] = \alpha$.

Let Y_j denote the binary random variable that takes value 1 if the total weight of balls thrown into the j -th bin is at least $\alpha \cdot 15/16$. Consequently, by Chernoff bound we have

$$\begin{aligned} \mathbb{E}[Y_j] &= \Pr[Y_j = 1] = 1 - \Pr\left[\sum_i X_j^i < \alpha \cdot \frac{15}{16}\right] \\ &\geq 1 - \exp\left\{-\frac{\alpha}{16^2 \cdot 2 \cdot \zeta}\right\}, \end{aligned} \tag{1}$$

which implies $\mathbb{E}[Y_j] > 127/128$ when ζ is small enough to promise that $\exp\{-0.01/(16^2 \cdot 2 \cdot \zeta)\} < 1/128$.

Let Z_j denote the binary random variable that takes value 1 if the total weight of balls thrown into the j -th bin is at most $\alpha \cdot 2$. Consequently, by Chernoff bound we have

$$\begin{aligned}\mathbb{E}[Z_j] &= \Pr[Z_j = 1] = 1 - \Pr\left[\sum_i X_j^i > \alpha \cdot 2\right] \\ &\geq 1 - \exp\left\{-\frac{\alpha}{3\zeta}\right\},\end{aligned}\tag{2}$$

which implies $\mathbb{E}[Z_j] > 127/128$ when ζ is small enough to promise that $\exp\{-0.01/(3\zeta)\} < 1/128$.

Hence, we have $\mathbb{E}[\sum_j Y_j] > H \cdot 127/128$ and $\mathbb{E}[\sum_j Z_j] > H \cdot 127/128$. Note that Y_1, \dots, Y_H are negatively associated, as well as Z_1, \dots, Z_H [7]. Hence, by Chernoff bound, it holds that with probability $1 - \exp\{-\Omega(H)\}$ there are at least $H \cdot 31/32$ bins with weight between $\alpha \cdot 15/16$ and $\alpha \cdot 2$. \square

Lemma 3. *Consider H bins and $l > H \cdot \Delta$ balls, where $\Delta > 2$. Balls are thrown to bins uniformly at random. If Δ is big enough, then with probability $1 - \exp\{-\Omega(H)\}$ there are at least $H \cdot 31/32$ bins that contain at least 2 balls.*

Proof. Considering the j -th bin, let X_j^i denote the random variable that takes value 1 if the i -th ball is in the j -th bin, and 0 otherwise. Since the balls are thrown to bins uniformly at random, then $\Pr[X_j^i = 1] = 1/F$ and $\mathbb{E}[\sum_i X_j^i] > \Delta$.

Let Y_j denote the binary random variable that takes value 1 if the number of balls thrown into the j -th bin is at least 2. Consequently, by Chernoff bound we have

$$\begin{aligned}\mathbb{E}[Y_j] &= \Pr[Y_j = 1] = 1 - \Pr\left[\sum_i X_j^i < 2\right] \\ &= 1 - \Pr\left[\sum_i X_j^i < \frac{2}{\Delta} \cdot \Delta\right] \\ &\geq 1 - \exp\left\{-\left(\frac{\Delta-2}{\Delta}\right)^2 \cdot \frac{\Delta}{2}\right\},\end{aligned}\tag{3}$$

which implies $\mathbb{E}[Y_j] > 63/64$ when Δ is big enough. By Chernoff bound, it holds that with probability $1 - \exp\{-\Omega(H)\}$ there are at least $H \cdot 31/32$ bins in which there are at least 2 balls. \square

Corollary 1. *Consider H bins and $l > H \cdot \Delta$ balls with weights $0 \leq w_1, w_2, \dots, w_l \leq \zeta$. Assume that $\sum_{i=1}^l w_i = \alpha \cdot H$ where $\alpha \geq 0.01$ is a constant. Balls are thrown to bins uniformly at random. Then, if Δ is big enough and ζ is small enough, then with probability $1 - \exp\{-\Omega(H)\}$ there are at least $H \cdot 15/16$ bins in which there are at least 2 balls, and the total weight is between $\alpha \cdot 15/16$ and $\alpha \cdot 2$.*

Proof. The conclusion is implied directly from Lemma 2 and Lemma 3. \square

At the end of this section, we introduce a result given in [3].

Lemma 4. *Consider a set of l nodes, v_1, v_2, \dots, v_l , transmitting on a channel. For node v_i , it transmits with probability $0 < p(v_i) < 1/2$. Let w_0 denote the probability that the channel is idle; and w_1 the probability that there is exactly one transmission on the channel. Then, $w_0 \cdot \sum_{i=1}^l p(v_i) \leq w_1 \leq 2 \cdot w_0 \cdot \sum_{i=1}^l p(v_i)$.*

The proof is omitted; readers can refer to [3] for the detailed proof.

3 Uniform Information Exchange

In this section, we introduce our Uniform Information Exchange (UIE) protocol. The pseudo-code of the protocol is given in Algorithm 1 and Algorithm 2.

UIE Protocol. There are two states for the nodes: *active* and *inactive*. Intuitively, the active nodes are trying to transmit messages over the network, while the inactive nodes just listen for incoming messages. Initially, all the source nodes are *active*, and the others are *inactive*.

In the protocol, an active node will become inactive when it successfully transmits its message to other active nodes. This way, on one hand, the number of active nodes is constantly decreasing, and on the other hand, it ensures that at any time the active nodes possess all k packets. Hence, when there is only one active

node left, it can send all the k packets to all the other nodes. The utilization of multiple channels can speed up the reduction of active nodes. By the transmissions on multiple channels, the active nodes can be reduced on all channels in parallel. However, as discussed before, when the number of active nodes becomes small, it cannot guarantee that for a particular channel, there are multiple active nodes operating on it. As a result, even if an active node successfully transmits on a channel, its message may not be received by other active nodes. In other words, the multiple channels are not efficient any more. Additionally, the protocol needs to ensure that when the surviving active node transmits, all other nodes listen on the same channel. Hence, we set a primary channel, which serves two purposes: first, it is used for reducing active nodes when the number of active nodes is small; second, it is used by the surviving active node to disseminate the packets.

Specifically, there are two processes in the protocol: the multiple-channel transmission process and the primary-channel transmission process. In the multiple-channel transmission process, active nodes operate on multiple channels to reduce the number of active nodes, while in the primary-channel transmission process, nodes operate on the primary channel. Note that because nodes have no idea about any network parameters, it is hard for nodes to determine when the multiple-channel transmission process should finish. Hence, in the protocol, these two processes are in parallel, rather than consecutive. Specifically, there are four slots in each round: in the first two slots, active nodes operate on multiple channels, and in the other slots, nodes operate on the primary channel. We set the first channel as the special primary channel. We next introduce the protocol in more detail.

Each active node v maintains two parameters $p(v)$ and $q(v)$. Denote the values of $p(v)$ and $q(v)$ in a round t by $p_t(v)$ and $q_t(v)$, respectively. In particular, $p_t(v)$ and $q_t(v)$ are the transmission probabilities of node v for the multi-channel transmission process and the primary-channel transmission process in round t , respectively. Initially, $p_0(v) := q_0(v) := \zeta$, where $0 < \zeta < 1$ is a constant (determined in Lemma 2). Let $m_t(v)$ denote the set of packets received by node v by round t . Initially, for a source node v initiated with packet P , $m_0(v) := \{P\}$. And for other nodes, $m_0(v) := \emptyset$.

The operations in the four slots of each round t are as follows:

- **Slot 1.** In this slot, the inactive nodes do nothing. Each active node v selects a channel from the \mathcal{F} candidates uniformly at random, and then transmits with probability $p_t(v)$ on the selected channel. If it does not transmit, it listens on the selected channel. If v receives a message containing a set of packets m' , it updates $m_{t+1}(v) := m' \cup m_t(v)$.

At the end of Slot 1, v updates the transmission probability p according to the following rule: if v listens and detects no transmission on the selected channel, $p_{t+1}(v) := \min\{\zeta, 2 \cdot p_t(v)\}$; otherwise, $p_{t+1}(v) := p_t(v)/2$.

- **Slot 2.** In this slot, the inactive nodes still do nothing. For an active node v , if it has received a message in Slot 1, it transmits an acknowledgement on the selected channel. Otherwise, v listens on the selected channel.

If an active node v transmitted in slot 1 and detects transmissions on the selected channel in Slot 2, the state of v switches to *inactive*.

- **Slot 3.** In this slot, all nodes operate on the primary channel (Channel 1). Specifically, all inactive nodes listen, and an active node v transmits with probability $q_t(v)$. At the end of Slot 3, active nodes update the transmission probability $q_t(\cdot)$ using the same rule as in Slot 1.

- **Slot 4.** For each (active or inactive) node v , if v received a message in Slot 3, it transmits an acknowledgement.

For an active node v , if v transmitted in Slot 3 and detects transmissions in this slot, it changes its state to *inactive*.

We state the correctness of the UIE protocol in the following Theorem 1.

Theorem 1. *Consider an execution of the UIE Protocol. When there is exactly one active node left, say node v in round T , then $p_T(v) = \bigcup_{v \in K} m_0(v)$. Recall that K is the set of all source nodes.*

Proof. Denote the set of active nodes in round t by A_t . Then $A_t \subseteq A_{t-1} \subseteq \dots \subseteq A_1 \subseteq A_0 = K$ holds for any $t > 0$, according to the protocol. Then the conclusion follows from the fact that $\bigcup_{v \in A_t} m_t(v) = \bigcup_{v \in A_{t-1}} m_{t-1}(v)$ holds for any $t > 0$, which is true because when an active node v becomes inactive in some round t , it means $m_t(v)$ is known to some other active node u which is still active in round $t + 1$. In detail, if an active node received acknowledgement or detected collisions in Slot 2, then it means its message has been received by some

Algorithm 1: UIE

Initialization: for node v at time 0

```
1  $p(v) := q(v) := \zeta$ ;  
2 if initially have a packet  $P$  then  
3    $m(v) := \{P\}$ ;  
4    $state(v) := Active$ ;  
5 else  
6    $m(v) := \{\}$ ;  
7    $state(v) := Inactive$ ;
```

Active State: for node v at time $t \geq 0$

```
8 Slot 1-2: pick a Channel  $r$  uniformly at random and call  $channel - use(r, m(v), p(v))$ ;  
9 Slot 3-4: call  $channel - use(1, m(v), q(v))$ ;
```

Inactive State: for node v at time $t \geq 0$

```
10 Slot 1-2: do nothing;  
11 Slot 3:  
12 listen on Channel 1;  
13 if receive a message containing a set of packets  $m'$  then  $m(v) := m(v) \cup m'$ ;  
14 Slot 4: if received a message in Slot 3 then transmit on Channel 1;
```

Algorithm 2: $channel - use(i, s, w)$

Slot 1:

```
1 on Channel  $i$ , transmit a message containing packets in  $s$  with probability  $w$  and listen with probability  $1 - w$ ;  
2 if listened then  
3   if Channel  $i$  is idle then  
4      $w := \min\{2w, \zeta\}$ ;  
5   else if received a message containing a set of packets  $m'$  then  
6      $s := s \cup m'$ ;  
7      $w := w/2$ ;  
8   else  
9     // Channel  $i$  is busy  
10     $w := w/2$ ;  
11 else  
12   // transmitted  
13    $w := w/2$ ;
```

Slot 2:

```
14 if received a message in Slot 1 then transmit on Channel  $i$ ;  
15 if transmitted in Slot 1 then  
16   listen on Channel  $i$ ;  
17   if receive a message OR Channel  $i$  is busy then  $state(v) := Inactive$ ;
```

other active nodes; if an active node received an acknowledgement or detected collisions in Slot 4, then it means its message has been received by all the other nodes in the network, including the active ones if any exists. \square

4 Analysis of the Protocol

In this section, we prove that with k source nodes, our protocol can disseminate all k packets to the whole network in $O(k/\mathcal{F} + \mathcal{F} \cdot \log n)$ rounds with high probability. Recall that \mathcal{F} is the number of available channels and n is the number of nodes in the network. Formally, this conclusion is summarized in Theorem 2.

Theorem 2. *Consider information exchange on a network of size n with \mathcal{F} available channels. For the case where there are initially $k \leq n$ source nodes, the following conclusions hold:*

1. *There exists a constant $\nu > 0$ such that with high probability, there is only one active node left at time $T^* := \nu(k/\mathcal{F} + \mathcal{F} \cdot \log n)$.*
2. *For time T^* when there is only one active node v left, at time $T^{**} := 2 \cdot T^* + \log n = O(T^*)$ it holds with high probability that every node in the network knows the k packets initially maintained by the source nodes and node v becomes inactive.*

Proof. The first conclusion follows directly from the Lemmas 5 and 13 which will be given later in Section 4.1 and 4.2, respectively. Here we first prove the second conclusion.

Since at time T^* there is only one active node v left, we know that $q_{2T^*}(v)$ will get back to ζ if node v is still active at time $2T^*$. Note that if node v transmits with probability ζ on the primary channel (Slot 3) for the subsequent $\log n$ rounds, then with high probability there exists one round in which node v transmits and consequently all the nodes in the network will receive the message. As shown in Theorem 1, the message transmitted by v contains all k packets. Hence, all nodes will get these packets in the received message. Finally, since the inactive nodes that received a message in Slot 3 transmit on the primary channel, then node v detects transmissions in Slot 4 and becomes inactive. \square

We next briefly introduce the analysis process for the first conclusion in Theorem 2. Recall that there are two parallel processes in our algorithm: the multi-channel transmission process (the first two slots in each round) and the primary-channel transmission process (the last two slots in each round). As discussed before, when there are many active nodes (more than $\mathcal{F} \cdot \log n$), multiple channels should be efficient in reducing the number of active nodes. When the number of active nodes is reduced to something small (less than $\mathcal{F} \cdot \log n$), the utilization of multiple channels might not be efficient any more, since for a particular channel, there might not be multiple nodes selecting it. In this case, we have to rely on the primary-channel transmission process to reduce the number of active nodes. Therefore, we divide the analysis into two parts. The first part analyzes how long it takes to decrease the number of active nodes to $\mathcal{F} \cdot \log n$ and the second part deals with how long it takes to further reduce the number of active nodes to one. More precisely, let A_t denote the set of active nodes in a round t . Let T be the first round in which the number of active nodes drops below $\mathcal{F} \cdot \log n$. That is, for any time $t < T$, it holds that $|A_t| \geq \mathcal{F} \cdot \log n$, and for any time $t \geq T$, it holds that $|A_t| < \mathcal{F} \cdot \log n$. Then the whole analysis is divided into two parts by T : The first part concerns the time period from 0 to $T - 1$, and the second part considers the algorithm execution since T . In the first part of the analysis, we mainly analyze the efficiency of the multi-channel transmission process in reducing the number of active nodes, and in the second part, we are mainly concerned about the efficiency of the primary-channel transmission process.

In the rest of this section, we assume that $k \geq \mathcal{F} \cdot \log n$. Otherwise, we can jump directly to the second part of the analysis.

4.1 Efficiency of Multiple Channels

In this section, we analyze the first part, i.e., the period from time 0 to the first round when the number of active nodes drops below $\mathcal{F} \cdot \log n$. The conclusion is summarized in the following Lemma 5.

Lemma 5. *There exists $T = O(k/\mathcal{F})$ such that in round T it holds with high probability that $|A_T| < \mathcal{F} \cdot \log n$.*

The main idea in proving Lemma 5 is to find a proper $\gamma' > 0$ such that after the protocol has been running for $T' = O(\log n)$ rounds, within any period of γ' rounds subsequently with $|A_t| \geq \mathcal{F} \cdot \log n$, there are (with constant probability) $\Omega(\mathcal{F})$ active nodes that switch from the *active* state to the *inactive* state. Then, with high

probability, there are $k - \mathcal{F} \cdot \log n < k$ active nodes switching to *inactive*, in a period of $O(\log n + k/\mathcal{F})$ (which is $O(k/\mathcal{F})$ for $k > \mathcal{F} \cdot \log n$) rounds. To prove Lemma 5, we need to introduce and prove a series of “small” lemmas at first, and leave the proof of Lemma 5 to the end of this section. We next do some preparation for proving Lemma 5.

By using more channels, it is natural to expect that the number of successful transmissions is increased accordingly. Specifically, it is expected that in a round, there should be $\Omega(\mathcal{F})$ successful transmissions with \mathcal{F} channels. In the following Lemma 6, we show that if a “safe range” on the total transmission probability of all active nodes is satisfied, the above expectation is true.

Lemma 6 (Safe range). *Consider the Uniform Information Exchange Protocol. For a round $t > 0$ with $|A_t| \geq \mathcal{F} \cdot \log n$, if there exist constants $\alpha_1, \alpha_2 \geq 1$ such that $\alpha_1 \cdot \mathcal{F} \leq \sum_{v \in A_t} p_t(v) \leq \alpha_2 \cdot \mathcal{F}$, then with constant probability there are $\Omega(\mathcal{F})$ active nodes switching to the inactive state in the second slot.*

Proof. For the convenience of the argument, we introduce a series of random variables $X^i(v)$ with $i = 1, \dots, \mathcal{F}$ and $v \in A_t$. The variable $X^i(v)$ takes value $p_t(v)$ if node v selects Channel i in the 1st slot of round t ; otherwise, $X^i(v) := 0$. Furthermore, denote $X^i := \sum_v X^i(v)$. By Corollary 1, with probability $1 - \exp\{-\Omega(\mathcal{F})\}$, there are at least $\mathcal{F} \cdot 15/16$ channels, such that for each of them there are at least two active nodes selecting it and the total transmission probability of these active nodes is between $\alpha_1 \cdot 15/16$ and $2 \cdot \alpha_2$. Next, we show that in such cases there are $\Omega(\mathcal{F})$ active nodes switching to the *inactive* state with constant probability.

With $b_i \in [0, 1/2]$ for $i = 0, 1, \dots$, it holds [6] that

$$4^{-\sum_i b_i} \leq \prod_i (1 - b_i) \leq e^{-\sum_i b_i}. \quad (4)$$

Hence, for Channel i with X^i between $\alpha_1 \cdot 15/16$ and $2 \cdot \alpha_2$, it is idle with probability at least $4^{-4\alpha_2}$, and there is exactly one transmission on the channel with probability at least $\alpha_1 \cdot 4^{-4\alpha_2} \cdot 15/16$ (by Lemma 4). If there are at least two active nodes selecting a channel and there is only one node transmitting on the channel, then the transmission will succeed and the one transmitting in the first slot will sense transmissions in the second slot. According to the algorithm, the node that transmitted will switch to the *inactive* state. Therefore there are at least $\mathcal{F} \cdot 15/16$ channels such that for each of them there is an active node switching to the *inactive* state with probability at least $\alpha_1 \cdot 4^{-4\alpha_2} \cdot 15/16$. In expectation, there are $C \cdot \mathcal{F}$ new *inactive* nodes where $C := (1 - \exp\{-\Omega(\mathcal{F})\})\alpha_1 \cdot 4^{-2\alpha_2} \cdot 15/16$ which is at least $\alpha_1 \cdot 4^{-2\alpha_2} \cdot 15/32$ when \mathcal{F} is large enough. Using the Chernoff bound over the \mathcal{F} channels, it holds that with constant probability (given α_1, α_2 , and large enough \mathcal{F}), there are $\Omega(\mathcal{F})$ active nodes switching to the *inactive* state in the second slot of round t . This completes the proof. \square

With the above Lemma 6, now the proof idea becomes clear, we only need to show that once initiated, the network will fall into the safe range very soon, and then stays in this range as long as there are enough active nodes, i.e. $|A_t| \geq \mathcal{F} \cdot \log n$. At the very beginning of the protocol, nodes are initiated with constant transmission probabilities, i.e. $p_0(\cdot) = \zeta$. Therefore, the summation of the initial transmission probabilities might be as large as $n \cdot \zeta$. We need to consider how long it takes for the summation $\sum_{v \in A_t} p_t(v)$ to drop below $\mathcal{F} \cdot \alpha_2$, where $\alpha_2 > 0$ is the constant defined by the safe range.

Lemma 7. *For a round t with $\sum_{v \in A_t} p_t(v) = \alpha \cdot \mathcal{F}$, it holds that $\Pr[\sum_{v \in A_t} p_{t+1}(v) \leq \alpha \cdot \mathcal{F} \cdot 3/4] \geq 7/8$ for large enough α .*

Proof. To show the conclusion, we need to look at some execution details of the UIE Protocol. Note that there are two parts concerning randomness. One part is in the channel selection, and the other part is in the transmission selection. Consider the channel selection part first, in which a random instance σ is a mapping from the $|A_t|$ active nodes to the \mathcal{F} channels. Recall that the probability of successful transmission on a channel is closely related to the total transmission probability of nodes selecting this channel. We call an instance *fair* if under it there are at least $\mathcal{F} \cdot 15/16$ channels such that on each of them the total transmission probability (of nodes selecting this channel) is at least $\alpha \cdot 15/16$. By Lemma 2, a fraction of $1 - \exp\{-\Omega(\mathcal{F})\}$ of instances are fair. We next consider such a fair instance σ .

Let X_σ be the random variable that indicates the value of $\sum_{v \in A_{t+1}} p_{t+1}(v)$, conditioned on channel selection instance σ . Clearly, $X_\sigma \leq 2 \sum_{v \in A_t} p_t(v)$, and for different instances σ , X_σ s are mutually independent. For a channel c , if without confusion, we also use c to denote the set of active nodes selecting channel c in the instance σ . Denote by X_σ^c the random variable that indicates the value of $\sum_{v \in c \cap A_{t+1}} p_{t+1}(v)$. Hence, $X_\sigma = \sum_c X_\sigma^c$.

Focus on a channel c with $\sum_{v \in c \cap A_t} p_t(v) \geq 15\alpha/16$. The probability that there is at least one transmission on channel c is at least $1 - \exp\{-15\alpha/16\}$, by Equation (4). According to the UIE Protocol, the nodes that selected channel c all halve their transmission probabilities if channel c is not idle in round t . Hence,

$$\begin{aligned} Pr[X_\sigma^c = \sum_{v \in c \cap A_t} \frac{p_t(v)}{2} \mid \sum_{v \in c \cap A_t} p_t(v) \geq \alpha \cdot \frac{15}{16}] \\ \geq 1 - \exp\{-\alpha \cdot \frac{15}{16}\}, \end{aligned}$$

which is at least $31/32$ when α is large enough. Hence in expectation, there are at least $(31/32) \cdot (\mathcal{F} \cdot 15/16)$ channels c with $X_\sigma^c = \sum_{v \in c \cap A_t} p_t(v)/2$.

Note that once the instance σ is given, the total transmission probability $\sum_{v \in c \cap A_t} p_t(v)$ for each channel c is specified. Then for different channels, the random variables X_σ^c s are mutually independent. Hence, by the Chernoff bound in Lemma 1, there are at least $(15/16) \cdot (\mathcal{F} \cdot 15/16)$ channels with $X_\sigma^c = \sum_{v \in c \cap A_t} p_t(v)/2$ with probability $1 - \exp\{-\Omega(\mathcal{F})\}$. Hence, with probability $1 - \exp\{-\Omega(\mathcal{F})\}$,

$$\begin{aligned} X_\sigma &\leq \mathcal{F} \cdot \left(\frac{15}{16}\right)^2 \cdot \frac{15\alpha}{16} \cdot \frac{1}{2} + (\alpha \cdot \mathcal{F} - \alpha \cdot \mathcal{F} \cdot \left(\frac{15}{16}\right)^3) \cdot 2 \\ &< \alpha \cdot \mathcal{F} \cdot 3/4. \end{aligned}$$

Finally it holds that $Pr[\sum_{v \in A_{t+1}} p_{t+1}(v) \leq \alpha \cdot \mathcal{F} \cdot 3/4] \geq (1 - \exp\{-\Omega(\mathcal{F})\}) \cdot (1 - \exp\{-\Omega(\mathcal{F})\})$ which is at least $7/8$ for large \mathcal{F} . The last thing to note is in the above analysis we did not consider the effect when an active node becomes *inactive*, which only makes the summation decrease and hence is not harmful. \square

Lemma 8 (Going down). *There exists a constant $\alpha'_2 > 1$, such that among $\gamma \log n$ rounds (not necessarily consecutive) with $\sum_{v \in A_t} p_t(v) \geq \alpha'_2 \cdot \mathcal{F}$ and sufficiently large $\gamma > 0$, there are at least $\frac{3}{4}\gamma \log n$ rounds with $\sum_{v \in A_{t+1}} p_{t+1}(v) < \frac{3}{4} \sum_{v \in A_t} p_t(v)$, with probability $1 - O(n^{-1})$.*

Proof. Let $T := \gamma \log n$, and X_t be the random variable that indicates the value of $\sum_{v \in A_{t+1}} p_{t+1}(v) / \sum_{v \in A_t} p_t(v)$. Then by Lemma 7, it holds that $Pr[X_t \leq 3/4] \geq 7/8$. Let Y_t be the binary random variable that takes value 1 if $X_t \leq 3/4$. Note that given $\sum_{v \in A_t} p_t(v) > \alpha'_2 \cdot \mathcal{F}$, $\mathbb{E}[Y_t] \geq 7/8$ always hold. Hence, $\mathbb{E}[\sum_{t=1}^T Y_t] \geq T \cdot 7/8$, and it holds that $Pr[\sum_{t=1}^T Y_t \leq T \cdot 3/4] = O(n^{-1})$ by the Chernoff bound. That is, with probability $1 - O(n^{-1})$, there are at least $T \cdot 3/4$ rounds t with $\sum_{v \in A_{t+1}} p_{t+1}(v) / \sum_{v \in A_t} p_t(v) \leq 3/4$, which completes the proof. \square

Lemma 9 (Fast adaptation). *There exists a constant $\alpha'_2 > 1$, such that during any period of $\gamma \log n$ rounds with sufficiently large $\gamma > 0$, the probability that within the considered period there is a round t with $\sum_{v \in A_t} p_t(v) \leq \alpha'_2 \cdot \mathcal{F}$ is $1 - O(n^{-1})$.*

Proof. Denote $T := \gamma \log n$. Without loss of generality, assume that the period of T rounds starts from $t = 1$ and ends at $t = T$, with $\sum_{v \in A_t} p_t(v) > \alpha'_2 \cdot \mathcal{F}$ always holds. Note that

$$\sum_{v \in A_T} p_T(v) = \sum_{v \in A_0} p_0(v) \cdot \prod_{t=0}^{T-1} \frac{\sum_{v \in A_{t+1}} p_{t+1}(v)}{\sum_{v \in A_t} p_t(v)}.$$

Then by Lemma 8, with probability at least $1 - O(n^{-1})$, it holds that

$$\begin{aligned} \sum_{v \in A_T} p_T(v) &\geq \sum_{v \in A_0} p_0(v) \cdot \left(\frac{3}{4} \cdot \frac{3}{4} \cdot \frac{3}{4} \cdot 2\right)^{\frac{T}{4}} \\ &= \sum_{v \in A_0} p_0(v) \cdot \left(\frac{27}{32}\right)^{\frac{T}{4}}, \end{aligned} \tag{5}$$

where the first inequality holds by coupling the “evolution” factors $\sum_{v \in A_{t+1}} p_{t+1}(v) / \sum_{v \in A_t} p_t(v)$. Since it holds that $\sum_{v \in A_0} p_0(v) < n$ and $T = \gamma \log n$, we know that $\sum_{v \in A_T} p_T(v)$ is at most $\alpha'_2 \cdot \mathcal{F}$ for large enough γ . \square

In the above, we have shown that the adaptation process of the total transmission probability (from the initial state to $\Theta(\mathcal{F})$) takes $O(\log n)$ rounds with high probability. Meanwhile, we also showed that when the total transmission probability increases beyond the upper bound of the safe range, the total transmission probability of active nodes shows a trend of going down. To finally show that in most of the rounds, the safe range is satisfied, we still need to show that if the total transmission probability of active nodes becomes very small, the trend is that it will go up.

Lemma 10. *There exists $\alpha' \geq 0.01$ such that for any time t with $\sum_{v \in A_t} p_t(v) = \alpha' \cdot \mathcal{F}$, it holds that*

$$\Pr\left[\sum_{v \in A_{t+1}} p_{t+1}(v) \geq \alpha' \cdot \mathcal{F} \cdot \frac{4}{3}\right] \geq \frac{7}{8}.$$

Proof. By Lemma 2, there is a fraction of $1 - \exp\{-\Omega(\mathcal{F})\}$ of channel selection instances in which on at least $\mathcal{F} \cdot 15/16$ channels the total transmission probability of active nodes selecting the channel is between $\alpha' \cdot 15/16$ and $\alpha' \cdot 2$ (i.e. the *fair* instances). Consider such an instance σ .

Let X_σ be the random variable that indicates the value of $\sum_{v \in A_{t+1}} p_{t+1}(v)$, conditioned on channel selection σ . Clearly, $X_\sigma \leq \sum_{v \in A_t} 2 \cdot p_t(v)$, and X_σ 's with distinct σ 's are independent. For a channel c , if without confusion, we also use c to denote the set of active nodes selecting c . Let X_σ^c denote the random variable that indicates the value of $\sum_{v \in c \cap A_{t+1}} p_{t+1}(v)$. Hence, $X_\sigma = \sum_c X_\sigma^c$. We consider a channel c with $\alpha' \cdot 15/16 \leq \sum_{v \in c \cap A_t} p_t(v) \leq \alpha' \cdot 2$. The probability that there are no transmissions on channel c is at least $4^{-2\alpha'}$. Hence $\Pr[X_\sigma^c = \sum_{v \in c \cap A_t} 2 \cdot p_t(v) | \alpha' \cdot 15/16 \leq \sum_{v \in c \cap A_t} p_t(v) \leq \alpha' \cdot 2] \geq 4^{-2\alpha'}$ which is at least $31/32$ when α' is close to 0.01. Hence, in expectation there are at least $(31/32) \cdot (\mathcal{F} \cdot 15/16)$ channels c with $X_\sigma^c = \sum_{v \in c \cap A_t} 2 \cdot p_t(v)$.

Note that once the instance σ is given, the total transmission probability of nodes $\sum_{v \in c \cap A_t} p_t(v)$ on each channel c is specified. Then for different channels, the random variables X_σ^c s are mutually independent. Hence, by the Chernoff bound, with probability $1 - \exp\{-\Omega(\mathcal{F})\}$ there are at least $(15/16) \cdot (\mathcal{F} \cdot 15/16)$ channels c with $X_\sigma^c = \sum_{v \in c \cap A_t} 2 \cdot p_t(v)$. Hence, with probability $1 - \exp\{-\Omega(\mathcal{F})\}$, $X_\sigma \geq (\mathcal{F} \cdot 15^2/16^2) \cdot (15\alpha'/16) \cdot 2 + (\alpha' \cdot \mathcal{F} - \alpha' \cdot \mathcal{F} \cdot 15^3/16^3)/2 - \zeta \cdot \mathcal{F} \cdot 17/16$, where the loss of weight $\zeta \cdot \mathcal{F} \cdot 17/16$ is due to those active nodes switching to the *inactive* state: at most $\mathcal{F}/16$ active nodes become *inactive* in the second slot, and at most $1 \leq \mathcal{F}$ active nodes become *inactive* in the 4th slot. When ζ is small enough, $\zeta \cdot \mathcal{F} \cdot 17/16$ is very small compared to $\mathcal{F} \cdot \alpha'$, and hence $X_\sigma \geq \mathcal{F} \cdot \alpha' \cdot 4/3$ (with probability $1 - \exp\{-\Omega(\mathcal{F})\}$). Finally it holds that $\Pr[\sum_{v \in A_{t+1}} p_{t+1}(v) \geq \alpha' \cdot \mathcal{F} \cdot 4/3] \geq (1 - \exp\{-\Omega(\mathcal{F})\}) \cdot (1 - \exp\{-\Omega(\mathcal{F})\})$ which is at least $7/8$ for large \mathcal{F} . \square

Lemma 11 (Going up). *There exists a constant $\alpha'_1 > 0$, such that among $\gamma \log n$ rounds (not necessarily consecutive) with $\sum_{v \in A_t} p_t(v) \leq \alpha'_1 \cdot \mathcal{F}$ and sufficiently large $\gamma > 0$, there are at least $\frac{3}{4}\gamma \log n$ rounds with $\sum_{v \in A_{t+1}} p_{t+1}(v) \geq \frac{4}{3} \sum_{v \in A_t} p_t(v)$, with probability $1 - O(n^{-1})$.*

Proof. Let $T := \gamma \log n$, and X_t be the random variable that indicates the value of $\sum_{v \in A_{t+1}} p_{t+1}(v) / \sum_{v \in A_t} p_t(v)$. Then by Lemma 10, it holds that $\Pr[X_t \geq 4/3] \geq 7/8$.

Let Y_t be the binary random variable that takes value 1 if $X_t \geq 2$. Note that given $\sum_v p_t(v) < \alpha'_1 \cdot \mathcal{F}$, $\mathbb{E}[Y_t] \geq 7/8$ always holds. Hence, $\mathbb{E}[\sum_{t=1}^T Y_t] \geq T \cdot 7/8$, and it holds that $\Pr[\sum_{t=1}^T Y_t \leq T \cdot 3/4] = O(n^{-1})$ by the Chernoff bound. That is, with probability $1 - O(n^{-1})$, there are at least $T \cdot 3/4$ rounds t with $\sum_{v \in A_{t+1}} p_{t+1}(v) / \sum_{v \in A_t} p_t(v) \geq 4/3$. \square

Now we are ready to show that in most of the rounds after the adaptation process, the total transmission probability of active nodes is in the safe range.

Lemma 12 (Stable). *Let t_0 be the first round in which $\sum_{v \in A_t} p_t(v)$ drops below $\alpha'_2 \cdot \mathcal{F}$. In the subsequent $T := \tau \cdot \log n$ rounds where $\tau > 0$ and n are large enough, the following hold:*

- (i) **hardly going high:** *there are at least $T \cdot 3/4$ rounds t with $\sum_{v \in A_t} p_t(v) \leq \alpha_2 \cdot \mathcal{F}$, where $\alpha_2 > \alpha'_2$ is a constant.*
- (ii) **hardly going low:** *there are at least $T \cdot 3/4$ rounds t with $\sum_{v \in A_t} p_t(v) \geq \alpha_1 \cdot k$, where $\alpha_1 < \alpha'_1$ is a constant.*

Proof. We prove the two conclusions one by one.

Proof for “hardly going high”. Consider the period from $t = t_0$ to $t = t_0 + T$. Define a *wave* to be an interval $[t_1, t_2]$ with $t_2 > t_1 + 19$, such that for rounds $t \in [t_1, t_2]$ it holds that $\sum_v p_t(v) > \alpha'_2 \cdot \mathcal{F}$, and for rounds

$t = t_1 - 1, t_2 + 1$ it holds that $\sum_{v \in A_t} p_t(v) \leq \alpha'_2 \cdot \mathcal{F}$. Then for any round t not in a wave, $\sum_{v \in A_t} p_t(v)$ is at most $\alpha_2 \cdot \mathcal{F}$ where $\alpha_2 := \alpha'_2 \cdot 2^{10}$.

Assume there are at least $T \cdot 1/4$ rounds t with $\sum_{v \in A_t} p_t(v) > \alpha_2 \cdot \mathcal{F}$. Otherwise, the lemma holds. Let \mathcal{A} denote the event that the assumption is true. Next, we show that \mathcal{A} will never happen when n is large enough. Let T' be the number of rounds t with $\sum_{v \in A_t} p_t(v) > \alpha_2 \cdot \mathcal{F}$. Clearly, these rounds are all on waves, and by the assumption, $T' \geq \frac{1}{4}T$. Let \mathcal{B} denote the event that among all these rounds, there are $T' \cdot 3/4$ rounds t with $X_t \leq 3/4$. Recall that X_t is the random variable that takes value $\sum_{v \in A_{t+1}} p_{t+1}(v) / \sum_{v \in A_t} p_t(v)$. Assume that $\tau > 4\gamma$, where γ is from Lemma 8. Then $T' > \gamma \log n$, and hence by Lemma 8, it holds that $\Pr[\mathcal{B}|\mathcal{A}] = 1 - O(n^{-1})$, which is positive when n is large enough. However, as shown in the following argument, events \mathcal{B} and \mathcal{A} do not happen together, which leads to the conclusion that \mathcal{A} will never happen when n is large enough.

Now we show that \mathcal{B} and \mathcal{A} do not happen together. Actually, it is sufficient to show that \mathcal{B} will not happen. Recall that event \mathcal{B} happens meaning that a fraction of $3/4$ rounds in waves satisfy $X_t \leq 3/4$. To show this is impossible, we focus on a single wave $[t_1, t_2]$, and prove that among these $t_2 - t_1 + 1$ rounds, there are less than $(t_2 - t_1 + 1) \cdot 3/4$ rounds t with $X_t \leq 3/4$. Assume the opposite, and then the value of $\sum_{v \in A_{t_2}} p_{t_2}(v)$ is at most $\sum_{v \in A_{t_1}} p_{t_1}(v) \cdot (27/32)^{(t_2 - t_1 + 1)/4}$ (using the coupling technique). Recalling that in a wave $t_2 - t_1 + 1 > 20$, we have $\sum_{v \in A_{t_2}} p_{t_2}(v) < \sum_{v \in A_{t_1}} p_{t_1}(v) \cdot (27/32)^5 < \sum_{v \in A_{t_1}} p_{t_1}(v)/2$. Since in round $t = t_1 - 1$, $\sum_{v \in A_t} p_t(v) < \alpha'_2 \cdot \mathcal{F}$, which implies that $\sum_{v \in A_{t_1}} p_{t_1}(v) \leq 2\alpha'_2 \cdot \mathcal{F}$. Hence, $\sum_{v \in A_{t_2}} p_{t_2}(v) < \alpha'_2 \cdot \mathcal{F}$, which contradicts the definition of the wave. Hence the assumption does not hold, which completes the proof.

Proof for “hardly going low”. Consider the period from $t = t_0$ to $t = t_0 + T$. Define a *hole* to be an interval $[t_1, t_2]$ with $t_2 > t_1 + 19$, such that for rounds $t = t_1, \dots, t_2$ it holds that $\sum_v p_t(v) < \alpha'_1 \cdot \mathcal{F}$, and for rounds $t = t_1 - 1, t_2 + 1$ it holds that $\sum_v p_t(v) \geq \alpha'_1 \cdot \mathcal{F}$. Then for any round t not in a hole, $\sum_v p_t(v)$ is at least $\alpha_1 \cdot k$ where $\alpha_1 := \alpha'_1/2^{10}$.

Assume there are at least $T \cdot 1/4$ rounds t with $\sum_{v \in A_t} p_t(v) < \alpha_1 \cdot \mathcal{F}$. Otherwise, the lemma holds. Let \mathcal{A} denote the event that the assumption is true. Next, we show that \mathcal{A} will never happen when n is large enough. Let T' be the number of rounds t with $\sum_v p_t(v) < \alpha_1 \cdot \mathcal{F}$. Clearly, these rounds are all in holes, and by the assumption, $T' \geq \frac{1}{4}T$. Let \mathcal{B} denote the event that among all these rounds, there are $T' \cdot 3/4$ rounds t with $X_t \geq 4/3$. Recall that X_t is the random variable that takes value $\sum_{v \in A_{t+1}} p_{t+1}(v) / \sum_{v \in A_t} p_t(v)$. Assume that $\tau > 4\gamma$, where γ is from Lemma 11. Then $T' > \gamma \log n$, and hence by Lemma 11, it holds $\Pr[\mathcal{B}|\mathcal{A}] = 1 - O(n^{-1})$, which is positive when n is large enough. However, as shown in the following argument, \mathcal{B} and \mathcal{A} do not happen together, which leads to the conclusion that \mathcal{A} will never happen when n is large enough.

Now we show that \mathcal{B} and \mathcal{A} never happen together. Actually, it is sufficient if we show \mathcal{B} never happen. Recall that if event \mathcal{B} happens, it means a fraction of $3/4$ of the considered rounds satisfy $X_t \geq 4/3$. To show this is impossible, we focus on a single hole $[t_1, t_2]$, and prove that among these $t_2 - t_1 + 1$ rounds, there are less than $(t_2 - t_1 + 1) \cdot 3/4$ rounds t with $X_t \geq 4/3$. Assume the opposite, and then the value of $\sum_{v \in A_{t_2}} p_{t_2}(v)$ is at least $\sum_{v \in A_{t_1}} p_{t_1}(v) \cdot (32/27)^{(t_2 - t_1 + 1)/4}$ (using the coupling technique). Recalling that in a hole $t_2 - t_1 + 1 > 20$, we have $\sum_{v \in A_{t_2}} p_{t_2}(v) > \sum_{v \in A_{t_1}} p_{t_1}(v) \cdot (32/27)^5 > \sum_{v \in A_{t_1}} 2 \cdot p_{t_1}(v)$. Since at $t = t_1 - 1$, $\sum_{v \in A_t} p_t(v) > \alpha'_1 \cdot \mathcal{F}$, which implies that $\sum_{v \in A_{t_1}} p_{t_1}(v) \geq \alpha'_1 \cdot \mathcal{F}/2$. Hence, $\sum_{v \in A_{t_2}} p_{t_2}(v) > \alpha'_1 \cdot \mathcal{F}$, which contradicts the definition of the hole. Hence, the hypothesis does not hold, which completes the proof. \square

Now, we are ready to prove Lemma 5.

Proof of Lemma 5. At first, recall that by Lemma 6, in any round t with $|A_t| \geq \mathcal{F} \cdot \log n$ and $\alpha_1 \cdot \mathcal{F} \leq \sum_{v \in A_t} p_t(v) \leq \alpha_2 \cdot \mathcal{F}$, there exists constants $0 < c_1, c_2 < 1$ such that with probability at least c_1 there are $c_2 \cdot \mathcal{F}$ active nodes switching to the *inactive* state.

Define T_1 as the first round t such that the summation $\sum_{v \in A_t} p_t(v)$ drops below $\alpha_2 \cdot k$. By Lemma 9, we know that $T_1 = O(\log n)$. After T_1 , by applying Lemma 12 it follows that for any period of length at least $T' := \max\{2 \cdot k/(\mathcal{F} \cdot c_1 \cdot c_2), \tau \cdot \log n\}$, with high probability, there are $T'/2$ rounds t in which $\sum_{v \in A_t} p_t(v)$ is between $\alpha_1 \cdot \mathcal{F}$ and $\alpha_2 \cdot \mathcal{F}$. Then we know that for large enough $\tau > 0$, with high probability there is a round $t < T_1 + T'$ that satisfies $|A_t| < \mathcal{F} \cdot \log n$. Otherwise, based on the above argument and using the Chernoff bound, it is easy to show that up to round $T_1 + T'$, there are more than k active nodes switching to the *inactive* state with high probability, which is impossible.

Hence, there exists constant $\gamma' > 0$ with $T := \gamma'(\log n + k/\mathcal{F}) \geq T_1 + T'$, such that with high probability there is a round $t \leq T$ that satisfies $|A_t| < \mathcal{F} \cdot \log n$. Recall that we assume $k \geq \mathcal{F} \cdot \log n$ (otherwise, we can ignore this section and only consider the analysis in Section 4.2), which implies $T = O(k/\mathcal{F})$. \square

4.2 Efficiency of the Primary Channel

In this section, we analyze the “second part” of the algorithm execution: the execution after the round when the number of active nodes drops below $\mathcal{F} \cdot \log n$. The conclusion is summarized in Lemma 13. Note that here in this part of the analysis, we do not consider the decrease of active nodes due to successful transmissions in the multi-channel transmission process. Since the multi-channel transmission process makes the decrease of active nodes much faster, the assumption will not affect the correctness of the analysis.

Lemma 13. *Consider a round T with $|A_T| \leq \mathcal{F} \cdot \log n$. There is a constant $\mu > 0$ such that at time $T^* \leq T + \mu \cdot \mathcal{F} \cdot \log n$ there is only one active node left with high probability.*

Proof. The proof for this lemma depends on a special case of the proof for Lemma 5, where $\mathcal{F} = 1$ and the transmission probability refers to $q(\cdot)$. Hence, we only give a brief sketch.

After time T with $|A_T| \leq \mathcal{F} \cdot \log n$, it takes at most $O(\log n)$ rounds for the summation $\sum_{v \in A_t} q_t(v)$ to fall down to a range between β_1 and β_2 . Here, β_1 and β_2 are constants such that for any round t with $\beta_1 \leq \sum_{v \in A_t} q_t(v) \leq \beta_2$, there is one active node switching to the *inactive* state in the 4th slot with constant probability. Afterward, consider a round $T' := T + \mu \cdot \mathcal{F} \cdot \log n$ where $\mu > 0$ is a large enough constant. Then with high probability there is a time round $t < T'$ such that $|A_t| = 1$. Otherwise, during the period from T to T' , with high probability there are more than $\mathcal{F} \cdot \log n$ active nodes switching to the *inactive* state in the 4th slot, which is impossible. \square

4.3 Stabilization

Recall in Lemma 9, we have proved that it takes $O(\log n)$ rounds for a network to become “safe”, which means the summation $\sum_{v \in A_t} p_t(v)$ goes from its initial value to a range between $\alpha_1 \cdot \mathcal{F}$ and $\alpha_2 \cdot \mathcal{F}$ for some constants $\alpha_1, \alpha_2 > 0$. This conclusion can be generalized to any network state that is not “safe”. We describe the generalized conclusion formally in the following Theorem 3.

Theorem 3. *Consider the case when the number of active nodes is always at least $\mathcal{F} \cdot \log n$. For a round t^* with $\sum_{v \in A_{t^*}} p_{t^*}(v)$ outside the safe range $[\alpha_1 \cdot \mathcal{F}, \alpha_2 \cdot \mathcal{F}]$, with high probability $\sum_{v \in A_t} p_t(v)$ will fall into the safe range in $\Phi = O(\log(\max\{\frac{p^*}{\mathcal{F}}, \frac{\mathcal{F}}{p^*}\}) + \log n)$ rounds.*

Proof. Recall that in the proof of Lemma 9, in order to show that the summation $\sum_{v \in A_t} p_t(v)$ goes below $\alpha_2 \cdot \mathcal{F}$, we considered $T := 4 \cdot T'$ rounds such that $T' \geq \gamma \log n$ for a large enough constant γ , during which there are $3 \cdot T'$ rounds with a decrease of $\sum_{v \in A_t} p_t(v)$ by a factor $3/4$ (by Lemma 8) and T' rounds with an increase of $\sum_{v \in A_t} p_t(v)$ by a factor at most 2. Then, after these T rounds, the summation $\sum_{v \in A_t} p_t(v)$ will be decreased by a factor of $(27/32)^{T'}$ with high probability. Since the network is initiated with $\sum_{v \in A_0} p_0(v) \leq \zeta \cdot n$, we know that it is enough to set $T := O(\log n)$ for the network to become “safe”.

In a similar approach, it is easy to show that for any round t^* with $p^* := \sum_{v \in A_{t^*}} p_{t^*}(v) > \alpha_2 \cdot \mathcal{F}$, by the round $t' := t^* + \max\{4 \cdot \log(32 \cdot p^* / (27 \cdot \alpha_2 \cdot \mathcal{F})), 4\gamma \log n\}$, the summation $\sum_{v \in A_{t'}} p_{t'}(v)$ becomes smaller than $\alpha_2 \cdot \mathcal{F}$ with high probability.

For the case that $p^* < \alpha_1 \cdot \mathcal{F}$, the proof idea is similar. Note that during $T := 4 \cdot T'$ rounds with $\sum_{v \in A_t} p_t(v) < \alpha_1 \cdot \mathcal{F}$, where $T' \geq \gamma \log n$ for a large enough constant γ , there are $3 \cdot T'$ rounds with an increase of $\sum_{v \in A_t} p_t(v)$ by a factor $4/3$ (Lemma 11) and T' rounds with a decrease of $\sum_{v \in A_t} p_t(v)$ by a factor $1/2$. Overall, after these T rounds, the summation $\sum_{v \in A_t} p_t(v)$ will be increased by a factor of $(32/27)^{T'}$ with high probability. Hence, by setting $t' := t^* + \max\{4 \cdot \log(27 \cdot \alpha_1 \cdot \mathcal{F} / (32 \cdot p^*)), 4\gamma \log n\}$, the summation $\sum_{v \in A_{t'}} p_{t'}(v)$ becomes larger than $\alpha_1 \cdot \mathcal{F}$ by round t' with high probability. \square

5 Conclusion

In this paper, we considered the information exchange problem of k source nodes in single-hop multiple-channel networks of n nodes. With \mathcal{F} available channels and collision detection, we proposed a protocol that solves the information exchange problem in $O(k/\mathcal{F} + \mathcal{F} \cdot \log n)$ rounds, with high probability. Our algorithm is uniform in n and k , which is the first known uniform algorithm for information exchange in multi-channel networks. And the proposed protocol is asymptotically optimal when k is large.

In our protocol, when detecting transmissions, a node will decrease its transmission probability to avoid collisions. Then if there exist jamming signals on a channel, an analysis similar to that introduced in this paper

would show that even for the case when jamming only affects a constant fraction of the available channels, the total transmission probability (i.e. $\sum_{v \in A_t} p_t(v)$) may tend to become very small. This affects the primary channel strategy even more significantly, since a fixed channel may be jammed all the time. This problem motivates us to consider jamming-resilience of the proposed protocol in the future.

References

- [1] A. F. Anta and M. A. Mosteiro. Contention resolution in multiple-access channels: k -selection in radio networks. In *Computing and Combinatorics, 16th Annual International Conference, COCOON 2010, Nha Trang, Vietnam, July 19-21, 2010. Proceedings*, pages 378–388, 2010.
- [2] A. F. Anta, M. A. Mosteiro, and J. R. Muñoz. Unbounded contention resolution in multiple-access channels. In *Proceedings of the 25th international conference on Distributed computing, DISC’11*, pages 225–236, Berlin, Heidelberg, 2011. Springer-Verlag.
- [3] B. Awerbuch, A. W. Richa, and C. Scheideler. A jamming-resistant MAC protocol for single-hop wireless networks. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Principles of Distributed Computing, PODC 2008, Toronto, Canada, August 18-21, 2008*, pages 45–54, 2008.
- [4] J. Capetanakis. Tree algorithms for packet broadcast channels. *IEEE Transactions on Information Theory*, 25(5):505–515, 1979.
- [5] B. S. Chlebus, G. D. Marco, and D. R. Kowalski. Scalable wake-up of multi-channel single-hop radio networks. In *Principles of Distributed Systems - 18th International Conference, OPODIS 2014, Cortina d’Ampezzo, Italy, December 16-19, 2014. Proceedings*, pages 186–201, 2014.
- [6] S. Daum, M. Ghaffari, S. Gilbert, F. Kuhn, and C. C. Newport. Maximal independent sets in multichannel radio networks. In *ACM Symposium on Principles of Distributed Computing, PODC ’13, Montreal, QC, Canada, July 22-24, 2013*, pages 335–344, 2013.
- [7] S. Daum, S. Gilbert, F. Kuhn, and C. Newport. Leader election in shared spectrum radio networks. In *Proceedings of the 2012 ACM symposium on Principles of distributed computing*, pages 215–224. ACM, 2012.
- [8] S. Daum, F. Kuhn, and C. C. Newport. Efficient symmetry breaking in multi-channel radio networks. In *Distributed Computing - 26th International Symposium, DISC 2012, Salvador, Brazil, October 16-18, 2012. Proceedings*, pages 238–252, 2012.
- [9] S. Dolev, S. Gilbert, R. Guerraoui, and C. C. Newport. Gossiping in a multi-channel radio network. In *Distributed Computing, 21st International Symposium, DISC 2007, Lemesos, Cyprus, September 24-26, 2007. Proceedings*, pages 208–222, 2007.
- [10] S. Dolev, S. Gilbert, R. Guerraoui, and C. C. Newport. Secure communication over radio channels. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Principles of Distributed Computing, PODC 2008, Toronto, Canada, August 18-21, 2008*, pages 105–114, 2008.
- [11] M. Ghaffari, S. Gilbert, C. C. Newport, and H. Tan. Optimal broadcast in shared spectrum radio networks. In *Principles of Distributed Systems, 16th International Conference, OPODIS 2012, Rome, Italy, December 18-20, 2012. Proceedings*, pages 181–195, 2012.
- [12] S. Gilbert, R. Guerraoui, D. R. Kowalski, and C. C. Newport. Interference-resilient information exchange. In *INFOCOM 2009. 28th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 19-25 April 2009, Rio de Janeiro, Brazil*, pages 2249–2257, 2009.
- [13] J. Hayes. An adaptive technique for local distribution. *Communications, IEEE Transactions on*, 26(8):1178–1186, Aug 1978.
- [14] S. Holzer, T. Locher, Y. A. Pignolet, and R. Wattenhofer. Deterministic multi-channel information exchange. In *24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA ’12, Pittsburgh, PA, USA, June 25-27, 2012*, pages 109–120, 2012.

- [15] S. Holzer, Y. A. Pignolet, J. Smula, and R. Wattenhofer. Time-optimal information exchange on multiple channels. In *FOMC'11, The Seventh ACM SIGACT/SIGMOBILE International Workshop on Foundations of Mobile Computing (part of FCRC 2011), San Jose, CA, USA, June 9, 2011, Proceedings*, pages 69–76, 2011.
- [16] D. R. Kowalski. On selection problem in radio networks. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing, PODC 2005, Las Vegas, NV, USA, July 17-20, 2005*, pages 158–166, 2005.
- [17] C. U. Martel. Maximum finding on a multiple access broadcast network. *Inf. Process. Lett.*, 52(1):7–15, 1994.
- [18] V. Mikhailov and B. Tsybackov. Free synchronous packet access in a broadcast channel with feedback. *Problemy Peredachi Inform*, 14(4), 2005.
- [19] W. Shi, Q. Hua, D. Yu, Y. Wang, and F. C. M. Lau. Efficient information exchange in single-hop multi-channel radio networks. In *Wireless Algorithms, Systems, and Applications - 7th International Conference, WASA 2012, Yellow Mountains, China, August 8-10, 2012. Proceedings*, pages 438–449, 2012.
- [20] Y. Wang, Y. Wang, D. Yu, J. Yu, and F. Lau. Information exchange with collision detection on multiple channels. *Journal of Combinatorial Optimization*, pages 1–18, 2014.
- [21] D. E. Willard. Log-logarithmic selection resolution protocols in a multiple access channel. *SIAM J. Comput.*, 15(2):468–477, May 1986.
- [22] Y. Yan, D. Yu, Y. Wang, J. Yu, and F. Lau. Bounded information dissemination in multi-channel wireless networks. *Journal of Combinatorial Optimization*, pages 1–17, 2014.